

Estimation of Covariance Matrices under Sparsity Constraints

Philippe Rigollet *

Alexandre B. Tsybakov †

March 3, 2013

Introduction. Estimation of covariance matrices in various norms is a critical issue that finds applications in a wide range of statistical problems, and especially in principal component analysis. It is well known that, without further assumptions, the empirical covariance matrix Σ^* is the best possible estimator in many ways, and in particular in a minimax sense. However, it is also well known that Σ^* is not an accurate estimator when the dimension p of the observations is high. The minimax analysis carried out by Tony Cai and Harry Zhou ([CZ] in what follows) guarantees that for several classes of matrices with reasonable structure (sparse or banded matrices), the fully data-driven thresholding estimator achieves the best possible rates when p is much larger than the sample size n . This is done, in particular, by proving minimax lower bounds that ensure that no estimator can perform better than the hard thresholding estimator, uniformly over the sparsity classes \mathcal{G}_q for each $0 \leq q < 1$. This result has a flavor of universality in the sense that one and the same estimator is minimax optimal for several classes of matrices.

Our comments focus on the sparsity classes of matrices.

- (a) *Optimal rates.* Optimal rates are obtained in [CZ] under the assumption that the dimension is very high: $p \geq n^\nu$, $\nu > 1$. Thus, the case of dimensions smaller than n , or even $p \approx n$, is excluded. This seems to be due to the technique of proving the lower bound (Theorem 2 in [CZ]). Indeed, by a different technique, we show that the lower bound holds without this assumption, cf. Theorem 1 below. Furthermore, in general, our lower rate $\psi^{(1)}$ is different from that obtained in [CZ] and has ingredients similar to the optimal rate for the Gaussian sequence model. We conjecture that it is optimal for all admissible configurations of n, p , and sparsity parameters.
- (b) *Frobenius norm and global sparsity.* We argue that the Frobenius norm is naturally adapted to the structure of the problem, at least for Gaussian observations, and we derive optimal rates under the Frobenius risk and global sparsity assumption.
- (c) *Approximate sparsity.* Again under the Frobenius risk, one can obtain not only the minimax results but also oracle inequalities. We demonstrate it for the soft-thresholding estimator. This allows us to deal with a more general setup where the covariance matrix is not necessarily sparse but can be well approximated by a sparse matrix.

*Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ 08544, USA. Supported in part by NSF grants DMS-0906424 and CAREER-DMS-1053987.

†Laboratoire de Statistique, CREST-ENSAE, 3, av. Pierre Larousse, F-92240 Malakoff Cedex, France. Supported in part by ANR under the grant “Parcimonie”.

Below we denote by $\|A\|$ the Frobenius norm of a matrix A :

$$\|A\|^2 = \text{Tr}(AA^\top) = \sum_{i,j} a_{ij}^2$$

where $\text{Tr}(B)$ stands for the trace of square matrix B . Moreover, for $q > 0$, we denote by $|v|_q$ the ℓ_q -norm of a vector v and by $|A|_q$ the ℓ_q norm of the off-diagonal entries of A . We set $|A|_0 = \sum_{i \neq j} I(a_{ij} \neq 0)$ (the number of non-zero off-diagonal entries of A). The operator $\ell_q \rightarrow \ell_q$ norm of A is denoted by $\|A\|_q$.

Frobenius norm and Sparsity. The cone of positive semi-definite (PSD) matrices can be equipped with a variety of norms, even more so than a vector space. [CZ] choose the $\|\cdot\|_1$ norm and consider classes of matrices that are essentially adapted to this metric. For example, the class \mathcal{G}_q defined in (1) controls the largest ℓ_q norm of the columns of the covariance matrix Σ with $0 \leq q < 1$ while the $\|\cdot\|_1$ norm measures the largest ℓ_1 norm of the columns of $\hat{\Sigma} - \Sigma$. Theorem 1 below indicates that for $q = 1$ consistent estimators do not exist.

One may wonder whether faster rates can be obtained if, for example, Σ has one row/column with large ℓ_q norm and all other rows/columns have small ℓ_q norm. It is quite clear that the $\|\cdot\|_1$ norm fails to capture such a behavior and we need to resort to other norms. As we see below, this is achievable when the Frobenius norm is used.

The Frobenius norm is a rather weak norm on the PSD cone. Indeed, it is very much a vector norm unlike the $\|\cdot\|_1$ norm used by [CZ] or the spectral norm that are operator norms. However, the choice of a norm is rather subjective but some general guidelines exist in a given statistical setup. It can be motivated by the idea of minimizing the Kullback-Leibler divergence between the true distribution and its estimator (see, e.g., [Rigollet, 2012](#)). This principle naturally gives rise to the use of the Frobenius norm in Gaussian covariance matrix estimation, as indicated by the following lemma.

Lemma 1. *Let I_p be the $p \times p$ identity matrix and Δ be a symmetric $p \times p$ matrix such that $I_p + \Delta$ is PSD. Denote by P_Σ the distribution of $\mathcal{N}_p(0, \Sigma)$ (a zero-mean normal random variable in \mathbb{R}^p with covariance matrix $\Sigma > 0$). Then, for any $0 < \varepsilon < 1$, the Kullback-Leibler divergence between $P_{I_p + \varepsilon \Delta}$ and P_{I_p} satisfies*

$$\text{KL}(P_{I_p + \varepsilon \Delta}, P_{I_p}) \leq \frac{g(-\varepsilon)}{2} \|\Delta\|^2,$$

where

$$g(\varepsilon) = \frac{\varepsilon - \log(1 + \varepsilon)}{\varepsilon^2}.$$

Moreover if $\|\Delta\|_2 \leq 1$, we have

$$\text{KL}(P_{I_p + \varepsilon \Delta}, P_{I_p}) \geq \frac{(1 - \log 2)\varepsilon^2}{2} \|\Delta\|^2. \quad (1)$$

PROOF. Take $\Sigma = I_p + \varepsilon \Delta$ and observe that

$$\text{KL}(P_\Sigma, P_{I_p}) = \mathbb{E} \log \left(\frac{dP_\Sigma}{dP_{I_p}}(X) \right) = \frac{1}{2} \mathbb{E} \log \left(\frac{1}{\det(\Sigma)} \right) + \frac{1}{2} \mathbb{E}[X^\top X - X^\top \Sigma^{-1} X],$$

where $X \sim \mathcal{N}_p(0, \Sigma)$. Let $\lambda_1, \dots, \lambda_p$ denote the eigenvalues of Δ and recall that $\det(\Sigma) = \prod_j (1 + \varepsilon \lambda_j)$. Moreover,

$$\mathbb{E}[X^\top X - X^\top \Sigma^{-1} X] = \mathbf{Tr}(\mathbb{E}[X X^\top] - \Sigma^{-1} \mathbb{E}[X X^\top]) = \mathbf{Tr}(\Sigma - I_p) = \sum_j \varepsilon \lambda_j.$$

Therefore,

$$\text{KL}(P_\Sigma, P_{I_p}) = \frac{1}{2} \sum_{j=1}^p [\varepsilon \lambda_j - \log(1 + \varepsilon \lambda_j)] \leq \frac{1}{2} \sum_{j=1}^p g(\varepsilon \lambda_j) \lambda_j^2.$$

Note now that since $I_p + \Delta$ is PSD, then $\lambda_j \geq -1$ for all $j = 1, \dots, p$. Therefore, since g is monotone decreasing on $(-1, \infty)$, it yields $g(\varepsilon \lambda_j) \leq g(-\varepsilon)$. The second statement of the lemma follows by observing that if $\|\Delta\|_2 \leq 1$, then $\varepsilon \lambda_j \leq \varepsilon \leq 1$ for all $j = 1, \dots, p$. \blacksquare

Minimax lower bounds over classes of sparse matrices. We denote by σ_{ij} the elements of Σ and by $\sigma_{(j)}$ the j th column of Σ with its j th component replaced by 0. For any $q > 0, R > 0$, we define the following classes of matrices:

$$\mathcal{G}_q^{(0)}(R) = \left\{ \Sigma \in \mathcal{C}_{>0} : |\Sigma|_q^q \leq R, \sigma_{ii} = 1, \forall i \right\},$$

$$\mathcal{G}_q^{(1)}(R) = \left\{ \Sigma \in \mathcal{C}_{>0} : \max_{1 \leq j \leq p} |\sigma_{(j)}|_q^q \leq R, \sigma_{ii} = 1, \forall i \right\},$$

where $\mathcal{C}_{>0}$ is the set of all positive definite symmetric $p \times p$ matrices. For $q = 0$, we define the classes $\mathcal{G}_0^{(0)}(R)$ and $\mathcal{G}_0^{(1)}(R)$ analogously, with the respective constraints $|\Sigma|_0 \leq R$ and $\max_{1 \leq j \leq p} |\sigma_{(j)}|_0 \leq R$. Here R is an integer for the class $\mathcal{G}_0^{(1)}(R)$, and an even integer for $\mathcal{G}_0^{(0)}(R)$ in view of the symmetry. We assume that $R = 2k \leq p(p-1)$ for $\mathcal{G}_0^{(0)}(R)$ and $R = k \leq p-1$ for $\mathcal{G}_0^{(1)}(R)$ where k is an integer. Set

$$\psi^{(0)} = R^{1/2} \left(\frac{1}{n} \log \left(1 + \frac{c_0 p^2}{R n^{q/2}} \right) \right)^{\frac{1}{2} - \frac{q}{4}}, \quad \psi^{(1)} = R \left(\frac{1}{n} \log \left(1 + \frac{c_0 p}{R n^{q/2}} \right) \right)^{\frac{1-q}{2}},$$

for some positive constant c_0 that does not depend on the parameters p, n, R . The following minimax lower bounds hold.

Theorem 1. Fix $R > 0, 0 \leq q \leq 2, C_0 > 0$ and integers $n \geq 1, p \geq 2$. Consider the conditions

$$R((\log p)/n)^{1-q/2} \leq C_0, \quad R((\log p)/n)^{(1-q)/2} \leq C_0, \quad R^{-1}((\log p)/n)^{q/2} \leq C_0. \quad (2)$$

Let X_1, \dots, X_n be i.i.d. $\mathcal{N}_p(0, \Sigma)$ random vectors, and let $w : [0, \infty) \rightarrow [0, \infty)$ be a monotone non-decreasing function such that $w(0) = 0$ and $w \not\equiv 0$. Then there exist constants $c_0 > 0, c_1 > 0, c > 0$ depending only on C_0 such that, under the first and third conditions in (2),

$$\inf_{\hat{\Sigma}} \sup_{\Sigma \in \mathcal{G}_q^{(0)}(R)} E_{\Sigma} w(\|\hat{\Sigma} - \Sigma\|/c_1 \psi^{(0)}) \geq c, \quad (3)$$

and under the second and third conditions in (2),

$$\inf_{\hat{\Sigma}} \sup_{\Sigma \in \mathcal{G}_q^{(1)}(R)} E_{\Sigma} w(\|\hat{\Sigma} - \Sigma\|_1 / c_1 \psi^{(1)}) \geq c, \quad \forall 0 \leq q \leq 1, \quad (4)$$

where E_{Σ} denotes the expectation with respect to the joint distribution of X_1, \dots, X_n and the infimum is taken over all estimators based on X_1, \dots, X_n .

PROOF. We first prove (3) with $q = 0$ and $R = 2k$. Assume first that $k \leq p^2/16$. We use Theorem 2.7 in Tsybakov (2009). It is enough to check that there exists a finite subset \mathcal{N} of $\mathcal{G}_0^{(0)}(2k)$ such that, for some constant $C > 0$ and some $\psi \geq C\psi^{(0)}$, we have

$$\begin{aligned} (i) \quad & \|\Sigma - \Sigma'\| \geq \psi, \quad \forall \Sigma \neq \Sigma' \in \mathcal{N} \cup \{I_p\}, \\ (ii) \quad & n \text{KL}(P_{\Sigma}, P_{I_p}) \leq 2^{-4} \log(\text{card } \mathcal{N}), \quad \forall \Sigma \in \mathcal{N}. \end{aligned}$$

We show that these conditions hold for

$$\psi = \left(\frac{k}{n} \log \left(1 + \frac{ep(p-1)}{2k} \right) \right)^{1/2}.$$

Let \mathcal{B} be the family of all $p \times p$ symmetric binary matrices, banded such that for all $B \in \mathcal{B}$, $b_{ij} = 0$ if $|i - j| > \sqrt{k}$, with 0 on the diagonal and exactly k nonzero over-diagonal entries equal to 1. Let M be the number of elements in the over-diagonal band where the entry 1 can only appear. For $k \leq p^2/4$ we have $M \geq p\sqrt{k} - k \geq p\sqrt{k}/2$. Therefore for, $k \leq p\sqrt{k}/4$, Lemma A.3 in Rigollet and Tsybakov (2011) implies that there exists a subset \mathcal{B}_0 of \mathcal{B} such that for any $B, B' \in \mathcal{B}_0, B \neq B'$, we have $\|B - B'\|^2 \geq (k+1)/4$, and

$$\log(\text{card } \mathcal{B}_0) \geq C_1 k \log \left(1 + \frac{ep}{4\sqrt{k}} \right) \quad (5)$$

for some absolute constant $C_1 > 0$. Consider the family of matrices $\mathcal{N} = \{\Sigma = I_p + \frac{a}{2}B : B \in \mathcal{B}_0\}$, where

$$a = a_0 \left(\frac{1}{n} \log \left(1 + \frac{ep}{4\sqrt{k}} \right) \right)^{1/2}$$

for some $a_0 > 0$. All matrices in \mathcal{N} have at most $2\sqrt{k}$ nonzero elements equal to a in each row. Therefore, the first inequality in (2) guarantees that for a_0 small enough, matrices $I_p + aB$ with $B \in \mathcal{B}_0$ and, a fortiori, $\Sigma \in \mathcal{N}$ are diagonally dominant and hence PSD. Thus, $\mathcal{N} \subset \mathcal{G}_0^{(0)}(2k)$ for sufficiently small $a_0 > 0$. Also, for any $\Sigma, \Sigma' \in \mathcal{N}, \Sigma \neq \Sigma'$, we have

$$\|\Sigma - \Sigma'\|^2 \geq C_2 a^2 k$$

for some absolute constant $C_2 > 0$. It is easy to see that this inequality also holds with a different C_2 if Σ or Σ' is equal to I_p . The above display implies (i). To check (ii), observe first that since $I_p + aB$ is PSD, we can apply Lemma 1 with $\Delta = aB, \varepsilon = 1/2$, to get

$$n \text{KL}(P_{\Sigma}, P_{I_p}) \leq \frac{na^2 g(-1/2)}{2} \|B\|^2 \leq a^2 kn, \quad \forall \Sigma \in \mathcal{N}.$$

To prove (ii), it suffices to take $a_0^2 < 2^{-4}C_1$, and to use (5). This proves (3) with $q = 0$ under the assumption $k \leq p^2/16$. The case $q = 0, k > p^2/16$ corresponds to a rate $\psi^{(0)}$ of order $\sqrt{p/n}$ and is easily treated via the Varshamov-Gilbert argument (we omit the details).

Next, observe that (3) for $0 < q \leq 2$, follows from the case $q = 0$. Indeed, let k be the maximal integer such that $2ka^q \leq R$ (we assume a_0 small enough to have $k \geq 1$, cf. the third inequality in (2)). Hence, $|\Sigma|_q^q = 2ka^q \leq R$ for any $\Sigma \in \mathcal{N}$. Also, $a\sqrt{k} \leq R^{1/2}a^{1-q/2}/\sqrt{2}$ and thus the first inequality in (2) ensures the positive definiteness of all $\Sigma \in \mathcal{N}$ for small a_0 . For this choice of k , we have $k + 1 > Ra^{-q}/2$ and $k \leq C_3 Rn^{q/2}$ with some constant $C_3 > 0$. It can be easily shown that (i) holds with

$$\psi^2 \geq CR \left(\frac{1}{n} \log \left(1 + \frac{ep^2}{Ra^{-q}} \right) \right)^{1-\frac{q}{2}} \geq CR \left(\frac{1}{n} \log \left(1 + \frac{c_0 p^2}{Rn^{q/2}} \right) \right)^{1-\frac{q}{2}}.$$

The proof of (4) is quite analogous, with the only difference that \mathcal{B} is now defined as the set of all symmetric binary matrices with exactly k off-diagonal entries equal to 1 in the first row and in the first column and all other entries 0. Then, for $k \leq (p-1)/2$, Lemma A.3 in [Rigollet and Tsybakov \(2011\)](#) implies that there exists a subset \mathcal{B}_1 of \mathcal{B} such that for any two distinct $B, B' \in \mathcal{B}_1$, we have $|b_{(1)} - b'_{(1)}|_1 \geq (k+1)/4$ (consequently, $\|B - B'\|_1 \geq (k+1)/4$) and

$$\log(\text{card } \mathcal{B}_1) \geq C_1 k \log \left(1 + \frac{e(p-1)}{k} \right). \quad (6)$$

Here, $b_{(1)}, b'_{(1)}$ are the first columns of B, B' with their first components replaced by 0. Thus, for any two distinct matrices Σ and Σ' belonging to the family $\mathcal{N}' = \{\Sigma = I_p + \frac{a}{2}B : B \in \mathcal{B}_1\}$ we have $\|\Sigma - \Sigma'\|_1^2 \geq C_4 a^2 k^2$ for some constant $C_4 > 0$. Here, $\mathcal{N}' \subset \mathcal{G}_0^{(1)}(k)$ thanks to the second inequality in (2). Also, by Lemma 1, $\text{KL}(P_\Sigma, P_{I_p}) \leq a^2 k$ for all $\Sigma \in \mathcal{N}'$. These remarks and (6) imply the suitably modified (i) and (ii) for the choice

$$a = a_0 \left(\frac{1}{n} \log \left(1 + \frac{e(p-1)}{k} \right) \right)^{1/2}$$

with a_0 small enough. The rest of the proof follows the same lines as the proof of (3). \blacksquare

The lower bound (4) and Theorem 4 in [CZ] imply that the rate $R((\log p)/n)^{\frac{1-q}{2}}$ is optimal on the class $\mathcal{G}_q^{(1)}(R)$ under the $\|\cdot\|_1$ norm if $Rn^{q/2} \leq p^\alpha$ with some $\alpha < 1$. In particular, for $q = 0$ this optimality holds under the quite natural condition $k = O(p^\alpha)$, and no lower bound on p in terms of n is required. Clearly, this is also true when we drop the condition $\Sigma > 0$ in the definition of $\mathcal{G}_q^{(1)}(R)$ and consider a weak ℓ_q constraint as in [CZ].

Note that the rate $\psi^{(0)}$ is very similar to the optimal rate in the Gaussian sequence model, cf. Section 11.5 in [Johnstone \(2011\)](#). This is due to the similarity between the vector ℓ_2 norm and the Frobenius norm. The rate $\psi^{(1)}$ is different but nevertheless has analogous ingredients. Observe also that, in contrast to the remark after Theorem 1 in [CZ], we prove the Frobenius and the $\|\cdot\|_1$ -norm lower bounds (3) and (4) by exactly the same technique. The key point is the use of the “ k -selection lemma” (Lemma A.3 in [Rigollet and Tsybakov \(2011\)](#)). The lower bound (4)

improves upon Theorem 2 in [CZ] in two aspects. First, it does not need the assumption $p \geq n^\nu$, $\nu > 1$, and provides insight on the presumed optimal rate for any configuration of n, p, R . Second, it is established for general loss functions w , in particular for the “in probability” loss that we consider below. The technique used in Theorem 2 of [CZ] is not adapted for this purpose as it applies to special losses derived from $w(t) = t$.

Approximate sparsity and optimal rates. Along with the hard thresholding estimator considered by [CZ], one can use the soft thresholding estimator $\tilde{\Sigma}$ defined as the matrix with off-diagonal elements

$$\tilde{\sigma}_{ij} = \text{sign}(\sigma_{ij}^*)(|\sigma_{ij}^*| - \tau)_+,$$

where σ_{ij}^* are the elements of the sample covariance matrix Σ^* , $\tau > 0$ is a threshold, and $(\cdot)_+$ denotes the positive part. The diagonal elements of $\tilde{\Sigma}$ are all set to 1 since we consider the classes $\mathcal{G}_q^{(j)}(R)$, $j = 0, 1$. Then $\tilde{\Sigma} = I_p + \tilde{\Sigma}_{\text{off}}$ where $\tilde{\Sigma}_{\text{off}}$ admits the representation (the minimum is taken over all $p \times p$ matrices S with zero diagonal):

$$\tilde{\Sigma}_{\text{off}} = \underset{S: \text{diag}(S)=0}{\text{argmin}} \left\{ |S - \Sigma^*|_2^2 + 2\tau|S|_1 \right\}.$$

Take the threshold

$$\tau = A\gamma\sqrt{\frac{\log p}{n}}, \quad (7)$$

where $A > 1$ and γ is the constant in the inequality (24) in [CZ].

Theorem 2. *Let X_1, \dots, X_n be i.i.d. random vectors in \mathbb{R}^p with covariance matrix Σ such that (24) in [CZ] holds. Assume that p, n , and A are such that $\tau \leq \delta$, where δ is the constant introduced after (24) in [CZ]. Then there exists $C_* > 0$ such that, with probability at least $1 - C_*p^{2-2A^2}$,*

$$\|\tilde{\Sigma} - \Sigma\|^2 \leq \min_S \left\{ \|S - \Sigma\|^2 + \left(\frac{1 + \sqrt{2}}{2} \right)^2 A^2 \gamma^2 \frac{|S|_0 \log p}{n} \right\}, \quad (8)$$

where \min_S denotes the minimum over all $p \times p$ matrices.

PROOF. Write $\sigma_{ij}^* = \sigma_{ij} + \xi_{ij}$ where the $\xi_{ij} = \sigma_{ij}^* - \sigma_{ij}$ are zero-mean random variables, $i \neq j$. Thus, considering σ_{ij}^* as observations, we have a sequence model in dimension $p(p-1)$. It is easy to see that it is a special case of the trace regression model studied in Koltchinskii et al. (2011) where A_0 is a diagonal matrix with the $p(p-1)$ off-diagonal entries of Σ on the diagonal. In the notation of Koltchinskii et al. (2011), the corresponding matrices X_i are diagonalizations of canonical basis vectors, the norm $\|\cdot\|_{L_2(\Pi)}$ coincides with the norm $|\cdot|_2$, and $\text{rank}(B)$ is equal to the number of non-zero entries of diagonal matrix B . Thus, Assumption 1 in Koltchinskii et al. (2011) is satisfied with $\mu = 1$, and we can apply Theorem 1 in Koltchinskii et al. (2011). It yields a deterministic statement:

$$|\tilde{\Sigma} - \Sigma|_2^2 \leq \min_S \left\{ |S - \Sigma|_2^2 + \left(\frac{1 + \sqrt{2}}{2} \right)^2 \tau^2 |S|_0 \right\}$$

provided $\tau > 2 \max_{i \neq j} |\sigma_{ij}^* - \sigma_{ij}|$. From (24) in [CZ] and a union bound, we obtain that, for τ defined in (7), this inequality holds with probability greater than $1 - C_* p^{2-2A^2}$. \blacksquare

Corollary 1. *Under the assumptions of Theorem 2, for any $0 < q < 2$, there exist constants $C', C_* > 0$ such that with probability at least $1 - C_* p^{2-2A^2}$,*

$$\|\tilde{\Sigma} - \Sigma\|^2 \leq \min_S \left\{ 2\|S - \Sigma\|^2 + C' |S|_q^q \left(\frac{\log p}{n} \right)^{1-\frac{q}{2}} \right\}. \quad (9)$$

PROOF. Let $|s_{[l]}|$, $l = 1, \dots, p(p-1)$, denote the absolute values of the off-diagonal elements of S ordered in a decreasing order. Note that for any $p \times p$ matrix S and any $0 < q < 2$ we have $|s_{[l]}|^q \leq |S|_q^q / l$. Fix an integer $k \leq p(p-1)$. Taking $s'_{ij} = s_{ij}$ if $|s_{ij}| \geq |s_{[k]}|$ and $s'_{ij} = 0$ otherwise, we get that for any S there exists a $p \times p$ matrix S' with $|S'|_0 = k$ such that

$$|S - S'|_2^2 = \sum_{l > k} s_{[l]}^2 \leq |S|_q^2 \sum_{l > k} l^{-2/q} \leq \frac{|S|_q^2 k^{1-2/q}}{2/q - 1}.$$

Together with Theorem 2, this implies that for any integer $k \leq p(p-1)$ we have

$$|\tilde{\Sigma} - \Sigma|_2^2 \leq \min_S \left\{ 2|S - \Sigma|_2^2 + \frac{|S|_q^2 k^{1-2/q}}{2/q - 1} + \left(\frac{1 + \sqrt{2}}{2} \right)^2 A^2 \gamma^2 \frac{k \log p}{n} \right\}.$$

Optimizing the right hand side over k completes the proof. \blacksquare

Note that the oracle inequalities (8) and (9) are satisfied for any covariance matrix Σ , not necessarily for sparse Σ . They quantify a trade-off between the approximation and sparsity terms. Their right-hand sides are small if Σ is well approximated by a matrix S with a small number of entries or with small ℓ_q norm of the off-diagonal elements. If the matrix Σ is sparse, $\Sigma \in \mathcal{G}_q^{(0)}(R)$, the oracle inequalities (8) and (9) imply that

$$\sup_{\Sigma \in \mathcal{G}_q^{(0)}(R)} P_\Sigma \left(\|\tilde{\Sigma} - \Sigma\| > C'' R^{1/2} \left(\frac{\log p}{n} \right)^{\frac{1}{2} - \frac{q}{4}} \right) \leq C_* p^{2-2A^2}$$

for some constant $C'' > 0$. This also holds when we drop the condition $\Sigma > 0$ in the definition of $\mathcal{G}_q^{(0)}(R)$. Combining this with Theorem 1, we find that the rate $R^{1/2} ((\log p)/n)^{\frac{1}{2} - \frac{q}{4}}$ is optimal on the class $\mathcal{G}_q^{(0)}(R)$ under the Frobenius norm if $R n^{q/2} \leq p^{2\alpha}$ with some $\alpha < 1$. In particular, for $q = 0$ this optimality holds under the condition $k \leq p^{2\alpha}$ with some $\alpha < 1$.

References

JOHNSTONE, I. M. (2011). Gaussian estimation: Sequence and wavelet models. Unpublished Manuscript., URL <http://www-stat.stanford.edu/~imj/>.

- KOLTCHINSKII, V., LOUNICI, K. and TSYBAKOV, A. (2011). Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *Ann. Statist.*, **39** 2302–2329.
- RIGOLLET, P. (2012). Kullback-Leibler aggregation and misspecified generalized linear models. *Ann. Statist. (to appear)*. *arXiv:0911.2919*.
- RIGOLLET, P. and TSYBAKOV, A. (2011). Exponential screening and optimal rates of sparse estimation. *The Annals of Statistics*, **39** 731–771.
- TSYBAKOV, A. B. (2009). *Introduction to nonparametric estimation*. Springer Series in Statistics, Springer, New York. Revised and extended from the 2004 French original, Translated by Vladimir Zaiats.